

Reflections for quantum query algorithms

Ben W. Reichardt

Abstract

We show that any boolean function can be evaluated optimally by a quantum query algorithm that alternates a certain fixed, input-independent reflection with a second reflection that coherently queries the input string. Originally introduced for solving the unstructured search problem, this two-reflections structure is therefore a universal feature of quantum algorithms.

Our proof goes via the general adversary bound, a semi-definite program (SDP) that lower-bounds the quantum query complexity of a function. By a quantum algorithm for evaluating span programs, this lower bound is known to be tight up to a sub-logarithmic factor. The extra factor comes from converting a continuous-time query algorithm into a discrete-query algorithm. We give a direct and simplified quantum algorithm based on the dual SDP, with a bounded-error query complexity that matches the general adversary bound.

Therefore, the general adversary lower bound is tight; it is in fact an SDP for quantum query complexity. This implies that the quantum query complexity of the composition $f \circ (g, \dots, g)$ of two boolean functions f and g matches the product of the query complexities of f and g , without a logarithmic factor for error reduction. It further shows that span programs are equivalent to quantum query algorithms.

1 Introduction

The query complexity, or decision-tree complexity, of a function measures the number of input bits that must be read in order to evaluate the function. Computation between queries is not counted. Quantum algorithms can run in superposition, and the quantum query complexity therefore allows coherent access to the input string. Quantum query complexity with bounded error lies below classical randomized query complexity, sometimes with a large gap [BV97, Sim97, Sho97, Aar09], but for total functions [BBC⁺01] or partial functions satisfying certain symmetries [AA09] the two measures are polynomially related; see the survey [BW02].

Although the query complexity of a function can fall well below its time complexity, studying query complexity has historically given insight into the power of quantum computers. For example, the quantum part of Shor's algorithms for integer factorization and discrete logarithm is a quantum query algorithm for period finding [Sho97]. Unlike for time complexity, there are also strong information-theoretic methods for placing lower bounds on quantum query complexity. These lower-bound techniques can be broadly classified as using either the polynomial method [BBC⁺01] or the adversary method [Amb02, ŠS06]. Høyer and Špalek [HŠ05] have surveyed the development of these two techniques and their multitude of applications. For now, suffice it to say that the two techniques are incomparable. In particular, for the n -input collision problem, the best adversary lower bound is of $O(1)$, whereas the correct complexity, determined by the polynomial method [AS04] and a matching algorithm [BHT98] is $\Theta(n^{1/3})$.

However, Høyer, Lee and Špalek [HLŠ07] discovered a strict generalization of the adversary bound that remains a lower bound on quantum query complexity:

Definition 1.1 (Adversary bounds). *For finite sets C and E , and $\mathcal{D} \subseteq C^n$, let $f : \mathcal{D} \rightarrow E$. An adversary matrix for f is a $|\mathcal{D}| \times |\mathcal{D}|$ real, symmetric matrix Γ that satisfies $\langle x | \Gamma | y \rangle = 0$ for all $x, y \in \mathcal{D}$ with $f(x) \neq f(y)$. Define the adversary and general adversary bounds for f by*

$$\text{Adv}(f) = \max_{\Gamma \geq 0} \frac{\|\Gamma\|}{\max_{j \in [n]} \|\Gamma \circ \Delta_j\|} \quad (1.1)$$

$$\text{Adv}^\pm(f) = \max_{\Gamma} \frac{\|\Gamma\|}{\max_{j \in [n]} \|\Gamma \circ \Delta_j\|} . \quad (1.2)$$

Both maximizations are over adversary matrices Γ , required to be entry-wise nonnegative in $\text{Adv}(f)$. $\Gamma \circ \Delta_j$ denotes the entry-wise matrix product between Γ and $\Delta_j = \sum_{x,y \in \mathcal{D}: x_j \neq y_j} |x\rangle\langle y|$.

Although the definitions of the two adversary bounds are very similar, the general adversary bound is much more powerful. In fact, the general adversary lower bound is always nearly tight:

Theorem 1.2 ([Rei09a]). *For any function $f : \mathcal{D} \rightarrow E$, with $\mathcal{D} \subseteq C^n$, the quantum query complexity $Q(f)$ satisfies*

$$Q(f) = O\left(\text{Adv}^\pm(f) \frac{\log \text{Adv}^\pm(f)}{\log \log \text{Adv}^\pm(f)} \log |C| \cdot \log |E|\right) . \quad (1.3)$$

This surprising upper bound follows from a connection between quantum query algorithms and the span program computational model [KW93] first observed in [RŠ08] and significantly strengthened in [Rei09a, Rei09b]. Note that the original statement [Rei09a, Theorem 10.2] of Theorem 1.2 restricted to the case $|C| = 2$ and included an additional factor of $\log \log |E|$ —this factor can be removed by [BNRW05, Corollary 3]. Lee has shown that the general adversary bound of a function with boolean output is stable under encoding the input into binary [Lee09], allowing the restriction $|C| = 2$ to be removed at a logarithmic cost.

Theorem 1.2 and the connection between span programs and quantum query algorithms behind its proof have corollaries including a query-optimal and nearly time-optimal quantum algorithm for evaluating a large class of read-once formulas over any finite gate set [Rei09c]. However, is Theorem 1.2 optimal? The factors of $\log |C|$ and $\log |E|$ are natural, but the log over log log term is not. It comes from converting a certain continuous-time query algorithm into a discrete-query algorithm following [CGM⁺09]. This conversion also somewhat obscures the algorithm’s structure.

It was conjectured that the unnatural log over log log factor could be removed [Rei09a, Conjecture 11.1]. In this article, we confirm the conjecture:

Theorem 1.3. *For any function $f : \mathcal{D} \rightarrow \{0, 1\}$, with $\mathcal{D} \subseteq \{0, 1\}^n$, the general adversary bound characterizes quantum query complexity:*

$$Q(f) = \Theta(\text{Adv}^\pm(f)) . \quad (1.4)$$

Theorem 1.3 suffices to simplify Eq. (1.3) following the proof of [Rei09a, Theorem 10.2]:

Corollary 1.4. *For finite sets C and E , $\mathcal{D} \subseteq C^n$, and any function $f : \mathcal{D} \rightarrow E$,*

$$Q(f) = \Omega(\text{Adv}^\pm(f)) \quad \text{and} \quad Q(f) = O(\text{Adv}^\pm(f) \log |C| \cdot \log |E|) . \quad (1.5)$$

Theorem 1.3 also allows for obtaining optimal results for the query complexity of composed functions. For $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g : \{0, 1\}^m \rightarrow \{0, 1\}$, let $f \bullet g$ be the function $\{0, 1\}^{nm} \rightarrow \{0, 1\}$ defined by $(f \bullet g)(x) = f(g(x_1, \dots, x_m), \dots, g(x_{(n-1)m+1}, \dots, x_{nm}))$. A bounded-error algorithm for evaluating $f \bullet g$ can be built from composing bounded-error algorithms for f and g ; thus, $Q(f \bullet g) = O(Q(f)Q(g) \log Q(f))$. However, the logarithmic factor for error reduction can be removed, and there is a matching lower bound:

Theorem 1.5. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g : \{0, 1\}^m \rightarrow \{0, 1\}$. Then*

$$Q(f \bullet g) = \Theta(Q(f)Q(g)) . \quad (1.6)$$

Proof. The general adversary bound composes as $\text{Adv}^\pm(f \bullet g) = \text{Adv}^\pm(f)\text{Adv}^\pm(g)$ for boolean functions f and g [HLS07, Rei09a], so the claim follows from Eq. (1.4). \square

The algorithm behind **Theorem 1.3** is substantially simpler than the algorithm for **Theorem 1.2**, although its analysis requires slightly more work. On input $x \in \{0, 1\}^n$, the algorithm consists of alternating applications of the input oracle O_x —a unitary that maps $|i, b\rangle$ to $|i, x_i \oplus b\rangle$, for $i = 1, \dots, n$ and $b \in \{0, 1\}$ —and a certain fixed reflection. The reflection is about the eigenvalue-zero subspace of the adjacency matrix A_G for a graph G derived from a dual SDP for Adv^\pm . This structure is based on Ambainis’s AND-OR formula-evaluation algorithm [Amb07]. A previous algorithm in Szegedy’s quantum walk model [Sze04] ran in $O(\text{Adv}^\pm(f) \|\text{abs}(A_G)\|)$ queries, where $\|\text{abs}(A_G)\|$ is the operator norm of the entry-wise absolute value of A_G [Rei09a, Prop. 9.4]. Ambainis’s approach efficiently removes the dependence on higher-energy portions of the adjacency matrix A_G . The analysis of the algorithm needs to transfer an “effective” spectral gap for the adjacency matrix of a related graph into an effective spectral gap for the product of O_x and the fixed reflection.

In fact, the input oracle is itself a reflection, $O_x^2 = \mathbf{1}$. Therefore the algorithm consists of alternating two fixed reflections, much like Grover’s search algorithm [Gro96]. It follows that every boolean function can be evaluated, with bounded error, optimally in this way. While known algorithms can in principle be converted into this form [Rei09a, Theorems 3.1, 5.2], we do not know an explicit closed form for the second reflection, e.g., for the collision problem.

The bipartite graph G can be thought of as a span program [KW93], and was constructed in the span program framework of [Rei09a]. Thus our algorithm is naturally seen as a quantum algorithm for evaluating span programs. Since the best span program for a function has witness size exactly equal to the general adversary bound [Rei09a, Rei10], **Theorem 1.3** also implies that quantum computers, measured by query complexity, and span programs, measured by witness size, are equivalent computational models for evaluating boolean functions. For simplicity, though, we will not detail this connection further. We will require from [Rei09a] only **Theorem 3.2** below, which without reference to span programs gives an effective spectral gap for a bipartite graph.

Barnum, Saks and Szegedy [BSS03] have given a family of SDPs that characterize quantum query complexity according to their feasibility or infeasibility, instead of according to the optimum value of a single SDP. The BSS SDPs work for any specified error rate, including zero. The general adversary bound is a polynomially smaller SDP, but of course the truth table of a function is typically exponentially long. Whereas our algorithm uses a workspace of $n + O(\log n)$ qubits to evaluate an n -bit boolean function (by [Rei09a, Lemma 6.6]), $n + 1$ qubits suffice by [BSS03]. To the author’s knowledge, neither the BSS SDPs nor Adv^\pm have ever been solved directly for a nontrivial, asymptotically large family of functions, with better than a constant-factor improvement over the adversary bound (see [CL08]). However, the easy composition rule for Adv^\pm , used in **Theorem 1.5**,

allows for computing Adv^\pm for a read-once formula by multiplying the bounds computed for constant-size gates. It may be that the very simple form of our algorithm will allow for further progress in the understanding and development of quantum query algorithms.

1.1 Open problems

An appealing open problem in quantum computing is to show a tighter relationship between classical and quantum query complexities for total functions—the largest known gap is quadratic but the best upper bound is $D(f) = O(Q(f)^6)$ [BBC⁺01]. Speculatively, the strong composition properties of quantum algorithms for total boolean functions may be a tool for approaching this problem. It also remains interesting to study non-boolean functions, their composition and the necessity of the $\log |C|$ and $\log |E|$ factors in Eq. (1.5). Theorem 1.5, the two-reflections form of the algorithm, and the elimination of the $\log \text{Adv}^\pm(f)$ factor suggest that it may be possible to adapt the algorithm to evaluate any boolean function f with a *bounded-error* input oracle with the same asymptotic number $\Theta(\text{Adv}^\pm(f))$ of quantum queries, following [HMW03] for the OR function. Classically, in the noisy decision-tree model, an extra logarithmic factor for error reduction is sometimes required [FRPU94], but this factor is not known to be needed for any quantum query algorithm [BNRW05].

1.2 Definitions

For a natural number $n \in \mathbf{N}$, let $[n] = \{1, 2, \dots, n\}$. For a bit $b \in \{0, 1\}$, let $\bar{b} = 1 - b$. For a finite set X , let \mathbf{C}^X be the Hilbert space $\mathbf{C}^{|X|}$ with orthonormal basis $\{|x\rangle : x \in X\}$. For vector spaces V and W over \mathbf{C} , let $\mathcal{L}(V, W)$ denote the set of all linear transformations from V into W , and let $\mathcal{L}(V) = \mathcal{L}(V, V)$. $\|A\|$ is the spectral norm of an operator A .

2 The algorithms

For a boolean function, taking the dual of the general adversary bound SDP in Definition 1.1 gives:

Lemma 2.1 ([Rei09a, Theorem 6.2]). *Let $f : \mathcal{D} \rightarrow \{0, 1\}$, with $\mathcal{D} \subseteq \{0, 1\}^n$. For $b \in \{0, 1\}$, let $F_b = \{x \in \mathcal{D} : f(x) = b\}$. Then*

$$\text{Adv}^\pm(f) = \min_{\substack{m \in \mathbf{N}, \{ |v_{xj}\rangle \in \mathbf{C}^m : x \in \mathcal{D}, j \in [n] \} : \\ \forall (x, y) \in F_0 \times F_1, \sum_{j \in [n] : x_j \neq y_j} \langle v_{xj} | v_{yj} \rangle = 1}} \max_{x \in \mathcal{D}} \sum_{j \in [n]} \| |v_{xj}\rangle \|^2 . \quad (2.1)$$

Based on a feasible solution to this SDP with objective value $W(\geq 1)$, we will give three algorithms for evaluating f , each with query complexity $O(W)$. (A feasible solution corresponds to a span program in canonical form, and its value equals the span program witness size [KW93, Rei09a].)

Let $I = [n] \times \{0, 1\} \times [m]$. Let $|t\rangle \in \mathbf{C}^{F_0}$ and $A \in \mathcal{L}(\mathbf{C}^{F_0}, \mathbf{C}^I)$ be given by

$$\begin{aligned} |t\rangle &= \frac{1}{3\sqrt{W}} \sum_{x \in F_0} |x\rangle \\ A &= \sum_{x \in F_0, j \in [n]} |x\rangle \langle j, \bar{x}_j| \otimes \langle v_{xj}| . \end{aligned} \quad (2.2)$$

Let G be the weighted bipartite graph with biadjacency matrix $B_G \in \mathcal{L}(\mathbf{C}^{\{0\}} \oplus \mathbf{C}^I, \mathbf{C}^{F_0})$:

$$B_G = \begin{pmatrix} |t\rangle & A \end{pmatrix} . \quad (2.3)$$

That is, G has a vertex for each row or column of B_G ; its vertex set is the disjoint union $F_0 \cup \{0\} \cup I$. Edges from F_0 to $\{0\} \cup I$ are weighted by the matrix entries. The weighted adjacency matrix of G is

$$A_G = \begin{pmatrix} 0 & B_G \\ B_G^\dagger & 0 \end{pmatrix} . \quad (2.4)$$

Let $\Delta \in \mathcal{L}(\mathbf{C}^{F_0 \cup \{0\} \cup I})$ be the orthogonal projection onto the span of all eigenvalue-zero eigenvectors of A_G . For an input $x \in \mathcal{D}$, let $\Pi_x \in \mathcal{L}(\mathbf{C}^{F_0 \cup \{0\} \cup I})$ be the projection

$$\Pi_x = \mathbf{1} - \sum_{j \in [n], k \in [m]} |j, \bar{x}_j, k\rangle \langle j, \bar{x}_j, k| . \quad (2.5)$$

That is, Π_x is a diagonal matrix that projects onto all vertices except those associated to the input bit complements \bar{x}_j . Finally, let

$$U_x = (2\Pi_x - \mathbf{1})(2\Delta - \mathbf{1}) . \quad (2.6)$$

U_x consists of the alternating reflections $2\Delta - \mathbf{1}$ and $2\Pi_x - \mathbf{1}$. The first reflection does not depend on the input x . The second reflection can be implemented using a single call to the input oracle O_x .

We present three related algorithms, each slightly simpler than the one before:

Algorithm 1:

1. Prepare the initial state $|0\rangle \in \mathbf{C}^{F_0 \cup \{0\} \cup I}$.
2. Run phase estimation on U_x , with precision $\delta_p = \frac{1}{100W}$ and error rate $\delta_e = \frac{1}{10}$.
3. Output 1 if the measured phase is zero. Otherwise output 0.

Algorithm 2:

1. Prepare the initial state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle \in \mathbf{C}^2 \otimes \mathbf{C}^{F_0 \cup \{0\} \cup I}$.
2. Pick $T \in [[100W]]$ uniformly at random. Apply the controlled unitary $|0\rangle\langle 0| \otimes \mathbf{1} + |1\rangle\langle 1| \otimes U_x^T$.
3. Measure the first register in the basis $\frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. Output 1 if the measurement result is $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and output 0 otherwise.

Algorithm 3:

1. Prepare the initial state $|0\rangle \in \mathbf{C}^{F_0 \cup \{0\} \cup I}$.
2. Pick $T \in [[10^5W]]$ uniformly at random. Apply U_x^T .
3. Measure the vertex. Output 1 if the measurement result is $|0\rangle$, and output 0 otherwise.

Phase estimation on a unitary V with precision δ_p and error rate δ_e can be implemented using $O(\frac{1}{\delta_p} \log \frac{1}{\delta_e})$ controlled applications of V [NWZ09], so the first algorithm has $O(W)$ query complexity. The second algorithm essentially applies a simplified version of phase estimation. Intuitively, it works because it suffices to distinguish zero from nonzero phase. The third algorithm does away with any phase estimation. Intuitively, this is possible because U_x is the product of two reflections, so its spectrum is symmetrical. The second and third algorithms clearly have $O(W)$ query complexity. The factor of 10^5 in the third algorithm's query complexity can be reduced by three orders of magnitude by adjusting downward the scaling factor for $|t\rangle$ in Eq. (2.2).

The time, or number of elementary operations, required to implement the reflection $2\Delta - \mathbf{1}$ is unclear. In practice it may still be preferable to use the potentially less query-efficient quantum walk algorithm from [Rei09a], as done for evaluating formulas in [Rei09c, RŠ08, Rei09d, ACR⁺10].

In the following section, we will show that all three algorithms correctly evaluate $f(x)$, with constant gaps between the soundness error and completeness parameters.

3 Analysis of the algorithms

To analyze the above algorithms, we shall study the spectrum of the unitary $U_x = (2\Pi_x - \mathbf{1})(2\Delta - \mathbf{1})$.

For this purpose, it will be useful to introduce two new graphs, following [RŠ08, Rei09a]. Let $\bar{\Pi}(x) = \sum_{j \in [n]} |j\rangle\langle j| \otimes |\bar{x}_j\rangle\langle \bar{x}_j| \otimes \mathbf{1}_{\mathbf{C}^I} \in \mathcal{L}(\mathbf{C}^I)$, and let $G(x)$ and $G'(x)$ be the weighted bipartite graphs with biadjacency matrices

$$B_{G(x)} = \begin{pmatrix} |t\rangle & A \\ 0 & \bar{\Pi}(x) \end{pmatrix} \quad \text{and} \quad B_{G'(x)} = \begin{pmatrix} A \\ \bar{\Pi}(x) \end{pmatrix}. \quad (3.1)$$

Based on the constraints of the SDP in Lemma 2.1, we can immediately construct eigenvalue-zero eigenvectors for $G(x)$ or $G'(x)$, depending on whether $f(x) = 1$ or $f(x) = 0$:

Lemma 3.1. *If $f(x) = 1$, let $|\psi\rangle = -3\sqrt{W}|0\rangle + \sum_{j \in [n]} |j, x_j\rangle \otimes |v_{x_j}\rangle \in \mathbf{C}^{\{0\} \cup I}$. Then $B_{G(x)}|\psi\rangle = 0$ and $|\langle 0|\psi\rangle|^2 / \|\psi\|^2 \geq 9/10$.*

If $f(x) = 0$, let $|\psi\rangle = -|x\rangle + \sum_{j \in [n]} |j, \bar{x}_j\rangle \otimes |v_{x_j}\rangle \in \mathbf{C}^{F_0 \cup I}$. Then $B_{G'(x)}^\dagger |\psi\rangle = 0$ and $|\langle t|\psi\rangle|^2 / \|\psi\|^2 \geq 1/(9W(W+1))$.

Let us recall from [Rei09a]:

Theorem 3.2 ([Rei09a, Theorem 8.7]). *Let G' be a weighted bipartite graph with biadjacency matrix $B_{G'} \in \mathcal{L}(\mathbf{C}^U, \mathbf{C}^T)$. Assume that $\delta > 0$ and $|t\rangle, |\psi\rangle \in \mathbf{C}^T$ satisfy $B_{G'}^\dagger |\psi\rangle = 0$ and $|\langle t|\psi\rangle|^2 \geq \delta \|\psi\|^2$.*

Let G be the same as G' except with a new vertex, 0, added to the U side, with outgoing edges weighted by the entries of $|t\rangle$. That is, the biadjacency matrix of G is

$$B_G = \begin{pmatrix} 0 & U \\ |t\rangle & B_{G'} \end{pmatrix} T \quad (3.2)$$

Let $\{|\alpha\rangle\}$ be a complete set of orthonormal eigenvectors of the weighted adjacency matrix A_G , with corresponding eigenvalues $\rho(\alpha)$. Then for all $\gamma \geq 0$, the squared length of the projection of $|0\rangle$ onto the span of the eigenvectors α with $|\rho(\alpha)| \leq \gamma$ satisfies

$$\sum_{\alpha: |\rho(\alpha)| \leq \gamma} |\langle \alpha|0\rangle|^2 \leq 8\gamma^2/\delta. \quad (3.3)$$

Substituting Lemma 3.1 into Theorem 3.2, we thus obtain the key statement:

Lemma 3.3. *If $f(x) = 1$, then $A_{G(x)}$ has an eigenvalue-zero eigenvector $|\psi\rangle$, supported on the column vertices, with*

$$\frac{|\langle 0|\psi\rangle|^2}{\|\psi\|^2} \geq \frac{9}{10} . \quad (3.4)$$

If $f(x) = 0$, let $\{|\alpha\rangle\}$ be a complete set of orthonormal eigenvectors of $A_{G(x)}$ with corresponding eigenvalues $\rho(\alpha)$. Then for any $c \geq 0$,

$$\sum_{\alpha: |\rho(\alpha)| \leq c/W} |\langle \alpha|0\rangle|^2 \leq 72 \left(1 + \frac{1}{W}\right) c^2 . \quad (3.5)$$

By choosing c a small positive constant, Eq. (3.5) gives an $O(1/W)$ “effective spectral gap” for eigenvectors of $A_{G(x)}$ supported on $|0\rangle$; it says that $|0\rangle$ has small squared overlap on the subspace of $O(1/W)$ -eigenvalue eigenvectors.

So far, we have merely repeated arguments from [Rei09a]. The main step in the analysis of our new algorithms is to translate Lemma 3.3 into analogous statements for U_x :

Lemma 3.4. *If $f(x) = 1$, then U_x has an eigenvalue-one eigenvector $|\varphi\rangle$ with*

$$\frac{|\langle 0|\varphi\rangle|^2}{\|\varphi\|^2} \geq \frac{9}{10} . \quad (3.6)$$

If $f(x) = 0$, let $\{|\beta\rangle\}$ be a complete set of orthonormal eigenvectors of U_x with corresponding eigenvalues $e^{i\theta(\beta)}$, $\theta(\beta) \in (-\pi, \pi]$. Then for any $\Theta \geq 0$,

$$\sum_{\beta: |\theta(\beta)| \leq \Theta} |\langle \beta|0\rangle|^2 \leq \left(2\sqrt{6\Theta W} + \frac{\Theta}{2}\right)^2 . \quad (3.7)$$

Assuming Lemma 3.4, the algorithms from Section 2 are both complete and sound. If $f(x) = 1$, then the first, phase-estimation-based algorithm outputs 1 with probability at least $9/10 - \delta_e = 4/5$. If $f(x) = 0$, then setting $\Theta = \delta_p = \frac{1}{100W}$, the algorithm outputs 1 with probability at most $\delta_e + (2\sqrt{6\Theta W} + \frac{\Theta}{2})^2 < 2/5$. The probability the second algorithm outputs 1 is the expectation versus T of $\frac{1}{4} \|(\mathbf{1} + U_x^T)|0\rangle\|^2$. If $f(x) = 1$, this is at least $9/10$ for all T . If $f(x) = 0$, let $\tau = \lceil 100W \rceil$ and simplify

$$\begin{aligned} \mathbb{E}_{T \in_R[\tau]} \left[\frac{1}{4} \|(\mathbf{1} + U_x^T)|0\rangle\|^2 \right] &= \mathbb{E}_{T \in_R[\tau]} \left[\frac{1}{4} \sum_{\beta} |1 + e^{i\theta(\beta)T}|^2 |\langle 0|\beta\rangle|^2 \right] \\ &= \frac{1}{4} \sum_{\beta} |\langle 0|\beta\rangle|^2 \left[2 + \frac{1}{\tau} \left(\frac{e^{i\theta(\beta)(\tau+1)} - e^{-i\theta(\beta)\tau}}{e^{i\theta(\beta)} - 1} - 1 \right) \right] . \end{aligned} \quad (3.8)$$

Setting $\Theta = \frac{1}{50W}$ and $\xi = (2\sqrt{6\Theta W} + \frac{\Theta}{2})^2$, we see that the algorithm outputs 1 with probability at most $\xi + (1 - \xi) \left(\frac{1}{2} + 1/(4\tau \sin \frac{\Theta}{2}) \right) < 88\%$ for $W \geq 1$. As its analysis requires more care, we defer consideration of the third algorithm to the end of this section.

For the proof of Lemma 3.4 we will use the following characterization of the eigen-decomposition of the product of reflections, essentially due to Jordan [Jor75]. Its use is common in quantum computation, e.g., [NWZ09, Sze04, MW05].

Lemma 3.5. *Given two projections Π and Δ , the Hilbert space can be decomposed into orthogonal one- and two-dimensional subspaces invariant under Π and Δ . On the one-dimensional invariant subspaces, $(2\Pi - \mathbf{1})(2\Delta - \mathbf{1})$ acts as either $+1$ or -1 . Each two-dimensional subspace is spanned by an eigenvalue- λ eigenvector $|v\rangle$ of $\Delta\Pi\Delta$, with $\lambda \in (0, 1)$, and $|v^\perp\rangle = (\mathbf{1} - \Delta)\Pi|v\rangle / \|(\mathbf{1} - \Delta)\Pi|v\rangle\|$. Letting $\theta = 2\arccos\sqrt{\lambda} \in (0, \pi)$, so $\Pi|v\rangle / \|\Pi|v\rangle\| = \cos\frac{\theta}{2}|v\rangle + \sin\frac{\theta}{2}|v^\perp\rangle$, the eigenvectors and corresponding eigenvalues of $(2\Pi - \mathbf{1})(2\Delta - \mathbf{1})$ on this subspace are, respectively,*

$$\frac{|v\rangle \mp i|v^\perp\rangle}{\sqrt{2}} \quad \text{and} \quad e^{\pm i\theta} . \quad (3.9)$$

Proof of Lemma 3.4. Notice from Eqs. (2.3) and (3.1) that G is naturally a subgraph of $G(x)$. Since $A_G\Delta = 0$ by definition of Δ , $A_{G(x)}\Delta = T(\mathbf{1} - \Pi_x)$, where T is a permutation matrix.

First consider the case $f(x) = 1$. Let $|\varphi\rangle$ be the restriction of $|\psi\rangle$ from Eq. (3.4) to the vertices of G . Since $|\psi\rangle$ has no support on the extra vertices of $G(x)$, $\|\varphi\| = \|\psi\|$ and $|\varphi\rangle$ is an eigenvalue-zero eigenvector of A_G ; $\Delta|\varphi\rangle = |\varphi\rangle$. Also $\Pi_x|\varphi\rangle = |\varphi\rangle$, so indeed $U_x|\varphi\rangle = |\varphi\rangle$.

Next consider the case $f(x) = 0$. Let

$$|\zeta\rangle = \sum_{\beta: |\theta(\beta)| \leq \Theta} |\beta\rangle \langle \beta|0\rangle \quad (3.10)$$

be the projection of $|0\rangle$ onto the space of eigenvectors with phase at most Θ in magnitude. Our aim is to upper bound $\|\zeta\|^2 = \langle 0|\zeta\rangle = |\langle 0|\hat{\zeta}\rangle|^2$, where $|\hat{\zeta}\rangle = |\zeta\rangle / \|\zeta\|$. Notice that $|\hat{\zeta}\rangle$ is supported only on eigenvectors $|\beta\rangle$ with $\theta(\beta) \neq 0$, i.e., on the two-dimensional invariant subspaces of Π_x and Δ . Indeed, if $U_x|\beta\rangle = |\beta\rangle$, then either $|\beta\rangle = \Pi_x|\beta\rangle = \Delta|\beta\rangle$ or $|\beta\rangle = (\mathbf{1} - \Pi_x)|\beta\rangle = (\mathbf{1} - \Delta)|\beta\rangle$. The first possibility implies $A_{G(x)}|\beta\rangle = 0$, so by Eq. (3.5) with $c = 0$, $\langle 0|\beta\rangle = 0$. In the second possibility, also $\langle 0|\beta\rangle = \langle 0|\Pi_x|\beta\rangle = 0$ since $|0\rangle = \Pi_x|0\rangle$.

We can split $\langle 0|\hat{\zeta}\rangle$ as

$$\begin{aligned} \langle 0|\hat{\zeta}\rangle &= \langle 0|\Delta|\hat{\zeta}\rangle + \langle 0|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle \\ &\leq |\langle 0|\Delta|\hat{\zeta}\rangle| + |\langle 0|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle| \\ &\leq |\langle 0|\Delta|\hat{\zeta}\rangle| + \|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle\| . \end{aligned} \quad (3.11)$$

Start by bounding the second term, $\|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle\|$. Intuitively, this term is small because $|\hat{\zeta}\rangle$ is supported only on two-dimensional invariant subspaces where Δ and Π_x are close. Indeed, let $|- \beta\rangle = (2\Delta - \mathbf{1})|\beta\rangle$, an eigenvector of A_G with phase $\theta(-\beta) = -\theta(\beta)$. Expanding $|\hat{\zeta}\rangle = \sum_\beta c_\beta |\beta\rangle$,

$$\begin{aligned} \|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle\|^2 &= \left\| \sum_\beta \Pi_x(\mathbf{1} - \Delta)c_\beta |\beta\rangle \right\|^2 \\ &= \sum_{\beta: \theta(\beta) > 0} \|\Pi_x(\mathbf{1} - \Delta)(c_\beta |\beta\rangle + c_{-\beta} |- \beta\rangle)\|^2 \\ &= \sum_{\beta: \theta(\beta) > 0} \sin^2 \frac{\theta(\beta)}{2} \|(\mathbf{1} - \Delta)(c_\beta |\beta\rangle + c_{-\beta} |- \beta\rangle)\|^2 \\ &\leq \left(\frac{\Theta}{2}\right)^2 \|(\mathbf{1} - \Delta)|\hat{\zeta}\rangle\|^2 . \end{aligned} \quad (3.12)$$

It remains to bound $|\langle 0|\Delta|\hat{\zeta}\rangle| = |\langle 0|w\rangle| \|\Delta|\hat{\zeta}\rangle\|$, where $|w\rangle = \Delta|\hat{\zeta}\rangle / \|\Delta|\hat{\zeta}\rangle\|$ is an eigenvalue-zero eigenvector of A_G . Intuitively, if $|\langle 0|w\rangle| = |\langle 0|\Pi_x|w\rangle|$ is large, then since A_G and $A_{G(x)}$ are the

same on Π_x , $\|A_{G(x)}|w\rangle\| = \|T(\mathbf{1} - \Pi_x)|w\rangle\|$ will be small. This in turn will imply that $|w\rangle$ has large support on the small-eigenvalue subspace of $A_{G(x)}$, contradicting Eq. (3.5).

Beginning the formal argument, we have

$$\begin{aligned}
\|A_{G(x)}\Delta|\hat{\zeta}\rangle\|^2 &= \|(\mathbf{1} - \Pi_x)\Delta|\hat{\zeta}\rangle\|^2 \\
&= \sum_{\beta:\theta(\beta)>0} \|(\mathbf{1} - \Pi_x)\Delta(c_\beta|\beta\rangle + c_{-\beta}|-\beta\rangle)\|^2 \\
&= \sum_{\beta:\theta(\beta)>0} \sin^2 \frac{\theta(\beta)}{2} \|\Delta(c_\beta|\beta\rangle + c_{-\beta}|-\beta\rangle)\|^2 \\
&\leq \left(\frac{\Theta}{2}\right)^2 \|\Delta|\hat{\zeta}\rangle\|^2.
\end{aligned} \tag{3.13}$$

Hence $\|A_{G(x)}|w\rangle\| \leq \Theta/2$.

Now split $|w\rangle = |w_{\text{small}}\rangle + |w_{\text{big}}\rangle$, where for a certain $d > 0$ to be determined,

$$\begin{aligned}
|w_{\text{small}}\rangle &= \sum_{\alpha:|\rho(\alpha)|\leq d\Theta/2} |\alpha\rangle\langle\alpha|w\rangle \\
|w_{\text{big}}\rangle &= \sum_{\alpha:|\rho(\alpha)|>d\Theta/2} |\alpha\rangle\langle\alpha|w\rangle.
\end{aligned} \tag{3.14}$$

Then

$$|\langle 0|\Delta|\hat{\zeta}\rangle| \leq |\langle 0|w\rangle| \leq |\langle 0|w_{\text{small}}\rangle| + |\langle 0|w_{\text{big}}\rangle|. \tag{3.15}$$

From Eq. (3.5) with $c = d\Theta W/2$, $|\langle 0|w_{\text{small}}\rangle| \leq \sqrt{72(1 + 1/W)}c\|w_{\text{small}}\rangle\| \leq 6d\Theta W$.

Since $A_{G(x)}|w\rangle = \sum_{\alpha} \rho(\alpha)|\alpha\rangle\langle\alpha|w\rangle$, we have

$$\begin{aligned}
\left(\frac{\Theta}{2}\right)^2 &\geq \|A_{G(x)}|w\rangle\|^2 \\
&= \|A_{G(x)}|w_{\text{small}}\rangle\|^2 + \|A_{G(x)}|w_{\text{big}}\rangle\|^2 \\
&\geq d^2 \left(\frac{\Theta}{2}\right)^2 \|w_{\text{big}}\rangle\|^2.
\end{aligned} \tag{3.16}$$

Hence $\|w_{\text{big}}\rangle\| \leq 1/d$.

Combining our calculations gives

$$\sqrt{\sum_{\beta:\theta(\beta)\leq\Theta} |\langle\beta|0\rangle|^2} = \langle 0|\hat{\zeta}\rangle \leq |\langle 0|\Delta|\hat{\zeta}\rangle| + \|\Pi_x(\mathbf{1} - \Delta)|\hat{\zeta}\rangle\| \leq 6d\Theta W + \frac{1}{d} + \frac{\Theta}{2}. \tag{3.17}$$

The right-hand side is $2\sqrt{6\Theta W} + \Theta/2$, as claimed, for $d = 1/\sqrt{6\Theta W}$. \square

Having proved Lemma 3.4, we return to the correctness proof for the third algorithm.

Proposition 3.6. *If $f(x) = 1$, then the third algorithm outputs 1 with probability at least 64%. If $f(x) = 0$, then the third algorithm outputs 1 with probability at most 61%.*

Proof. Letting $\tau = \lceil 10^5 W \rceil$, the third algorithm outputs 1 with probability

$$p_1 := \mathbb{E}_{T \in R[\tau]} [\langle 0 | U_x^T | 0 \rangle^2] = \mathbb{E}_{T \in R[\tau]} \left| \sum_{\beta} e^{i\theta(\beta)T} |\langle \beta | 0 \rangle|^2 \right|^2. \quad (3.18)$$

If $f(x) = 1$, then a crude bound puts p_1 at least $(9/10 - 1/10)^2 = 64\%$.

Assume $f(x) = 0$. Recall the notation that for an eigenvector $|\beta\rangle$ with $|\theta(\beta)| \in (0, \pi)$, $|\beta\rangle = (2\Delta - \mathbf{1})|\beta\rangle$ denotes the corresponding eigenvector with eigenvalue phase $\theta(-\beta) = -\theta(\beta)$. The key observation for this proof is that

$$\langle 0 | \beta \rangle = e^{-i\theta(\beta)} \langle 0 | -\beta \rangle. \quad (3.19)$$

This equal splitting of $|\langle 0 | \beta \rangle|$ and $|\langle 0 | -\beta \rangle|$ will allow us to bound p_1 close to $1/2$ instead of the trivial bound $p_1 \leq 1$. The intuition is that after applying U_x a suitable number of times, eigenvectors $|\beta\rangle$ and $|\beta\rangle$ will accumulate roughly opposite phases, so their overlaps with $|0\rangle$ will roughly cancel out. For this argument to work, though, the eigenvalue phase $\theta(\beta)$ should be bounded away from zero and from $\pm\pi$. Therefore define the projections

$$\begin{aligned} \Delta_{\Theta} &= \sum_{\beta: |\theta(\beta)| \leq \Theta} |\beta\rangle\langle\beta| \\ \overline{\Delta}_{\Lambda} &= \sum_{\beta: |\theta(\beta)| > \Lambda} |\beta\rangle\langle\beta| \\ \Sigma &= \mathbf{1} - \Delta_{\Theta} - \overline{\Delta}_{\Lambda}, \end{aligned} \quad (3.20)$$

where Θ and Λ , $0 < \Theta < \Lambda < \pi$, will be determined below. [Lemma 3.4](#) immediately gives the bound $\|\Delta_{\Theta}|0\rangle\| \leq 2\sqrt{6\Theta W} + \frac{\Theta}{2}$. We can also place a bound on $\|\overline{\Delta}_{\Lambda}|0\rangle\|$, using

$$2(\Delta - \mathbf{1})|0\rangle = (U_x^{\dagger} - \mathbf{1})|0\rangle = \sum_{\beta} (e^{-i\theta(\beta)} - 1)|\beta\rangle\langle\beta|0\rangle. \quad (3.21)$$

Expanding the squared norm of both sides gives

$$\|(U_x^{\dagger} - \mathbf{1})|0\rangle\|^2 = 4 \sum_{\beta} \sin^2 \frac{\theta(\beta)}{2} |\langle \beta | 0 \rangle|^2 \geq \|\overline{\Delta}_{\Lambda}|0\rangle\|^2 \cdot 4 \sin^2 \frac{\Lambda}{2} \quad (3.22)$$

and

$$\|(U_x^{\dagger} - \mathbf{1})|0\rangle\|^2 = 4(1 - \|\Delta|0\rangle\|^2) \leq 2/5. \quad (3.23)$$

In the second step we have used that $\|\Delta|0\rangle\|^2 \geq 9/10$; provided that f is not the constant zero function, A_G must have an eigenvalue-zero eigenvector with large overlap on $|0\rangle$. Combining Eqs. (3.22) and (3.23) gives

$$\|\overline{\Delta}_{\Lambda}|0\rangle\|^2 \leq \frac{1}{10 \sin^2 \frac{\Lambda}{2}}. \quad (3.24)$$

Returning to Eq. (3.18), we have

$$\begin{aligned} p_1 &\leq \mathbb{E}_{T \in R[\tau]} \left(\|\Delta_{\Theta}|0\rangle\|^2 + \|\overline{\Delta}_{\Lambda}|0\rangle\|^2 + \left| \sum_{\beta: \theta(\beta) \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 (e^{i\theta(\beta)T} + e^{-i\theta(\beta)T}) \right|^2 \right) \\ &\leq (\|\Delta_{\Theta}|0\rangle\|^2 + \|\overline{\Delta}_{\Lambda}|0\rangle\|^2) (2 - \|\Delta_{\Theta}|0\rangle\|^2 - \|\overline{\Delta}_{\Lambda}|0\rangle\|^2) \\ &\quad + \mathbb{E}_{T \in R[\tau]} \left(\sum_{\beta: \theta(\beta) \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 (e^{i\theta(\beta)T} + e^{-i\theta(\beta)T}) \right)^2. \end{aligned} \quad (3.25)$$

The algorithm chooses T at random to allow bounding the last term. Expanding this term gives

$$\begin{aligned}
& \mathbb{E}_{T \in_R[\tau]} \sum_{\beta, \beta': \theta(\beta), \theta(\beta') \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 |\langle \beta' | 0 \rangle|^2 (e^{i\theta(\beta)T} + e^{-i\theta(\beta)T}) (e^{i\theta(\beta')T} + e^{-i\theta(\beta')T}) \\
&= \mathbb{E} \sum_{\theta, \theta' \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 |\langle \beta' | 0 \rangle|^2 ((e^{i(\theta+\theta')T} + e^{-i(\theta+\theta')T}) + (e^{i(\theta-\theta')T} + e^{-i(\theta-\theta')T})) \\
&\leq \frac{1}{2} \|\Sigma|0\rangle\|^4 + \mathbb{E} \sum_{\theta, \theta' \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 |\langle \beta' | 0 \rangle|^2 (e^{i(\theta+\theta')T} + e^{-i(\theta+\theta')T}) \\
&= \frac{1}{2} \|\Sigma|0\rangle\|^4 + \frac{1}{\tau} \sum_{\theta, \theta' \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 |\langle \beta' | 0 \rangle|^2 \left(\frac{e^{i(\theta+\theta')(\tau+1)} - e^{-i(\theta+\theta')\tau}}{e^{i(\theta+\theta')} - 1} - 1 \right) \\
&\leq \frac{1}{2} \left(1 + \frac{1/\tau}{\min_{\theta, \theta' \in (\Theta, \Lambda]} |e^{i(\theta+\theta')} - 1|} \right) \|\Sigma|0\rangle\|^4 \\
&\leq \frac{1}{2} \left(1 + \frac{1}{2\tau \min\{\sin \Theta, \sin \Lambda\}} \right) \|\Sigma|0\rangle\|^4.
\end{aligned} \tag{3.26}$$

Here for brevity we have written θ and θ' for $\theta(\beta)$ and $\theta(\beta')$, respectively. In the second and fourth steps, we have used $\sum_{\theta \in (\Theta, \Lambda]} |\langle \beta | 0 \rangle|^2 = \frac{1}{2} \|\Sigma|0\rangle\|^2$. In the last step, we have used $|e^{i(\theta+\theta')} - 1| = 2 \sin \frac{\theta+\theta'}{2} \geq 2 \min\{\sin \Theta, \sin \Lambda\}$. Substituting the result back into Eq. (3.25) gives

$$\begin{aligned}
p_1 &\leq 1 - \frac{1}{2} \left(1 - \frac{1}{2\tau \min\{\sin \Theta, \sin \Lambda\}} \right) \|\Sigma|0\rangle\|^4 \\
&\leq 1 - \frac{1}{2} \left(1 - \frac{1}{2\tau \min\{\sin \Theta, \sin \Lambda\}} \right) \max \left[1 - \frac{1}{10 \sin^2 \frac{\Lambda}{2}} - \left(2\sqrt{6\Theta W} + \frac{\Theta}{2} \right)^2, 0 \right]^2.
\end{aligned} \tag{3.27}$$

Setting $\Lambda = \pi - \Theta$ and $\Theta = 1/(2000W)$, for $W \geq 1$ a calculation yields $p_1 \leq 61\%$. \square

Acknowledgements

I would like to thank Troy Lee for his help in formulating [Theorem 1.2](#). I also thank Sergio Boixo, Stephen Jordan, Julia Kempe and Rajat Mittal for helpful comments, and the Institute for Quantum Information for hospitality. Research supported by NSERC, ARO and MITACS.

References

- [AA09] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. 2009, [arXiv:0911.0996](#) [[quant-ph](#)].
- [Aar09] Scott Aaronson. BQP and the polynomial hierarchy. 2009, [arXiv:0910.4698](#) [[quant-ph](#)].
- [ACR⁺10] Andris Ambainis, Andrew M. Childs, Ben W. Reichardt, Robert Špalek, and Shengyu Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM J. Comput.*, 39(6):2513–2530, 2010. Earlier version in FOCS'07.
- [Amb02] Andris Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. Syst. Sci.*, 64:750–767, 2002, [arXiv:quant-ph/0002066](#). Earlier version in STOC'00.

- [Amb07] Andris Ambainis. A nearly optimal discrete query quantum algorithm for evaluating NAND formulas. 2007, [arXiv:0704.3628 \[quant-ph\]](#).
- [AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problem. *J. ACM*, 51(4):595–605, 2004.
- [BBC⁺01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001, [arXiv:quant-ph/9802049](#). Earlier version in FOCS’98.
- [BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum algorithm for the collision problem. In *Proc. 3rd LATIN*, LNCS vol. 1380, pages 163–169, 1998, [arXiv:quant-ph/9705002](#).
- [BNRW05] Harry Buhrman, Ilan Newman, Hein Röhrig, and Ronald de Wolf. Robust polynomials and quantum algorithms. In *Proc. 22nd STACS*, LNCS vol. 3404, pages 593–604, 2005, [arXiv:quant-ph/0309220](#).
- [BSS03] Howard Barnum, Michael Saks, and Mario Szegedy. Quantum query complexity and semidefinite programming. In *Proc. 18th IEEE Complexity*, pages 179–193, 2003.
- [BV97] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997. Earlier version in STOC’93.
- [BW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002.
- [CGM⁺09] Richard Cleve, Daniel Gottesman, Michele Mosca, Rolando Somma, and David L. Yonge-Mallo. Efficient discrete-time simulations of continuous-time quantum query algorithms. In *Proc. 41st ACM STOC*, pages 409–416, 2009, [arXiv:0811.4428 \[quant-ph\]](#).
- [CL08] Andrew M. Childs and Troy Lee. Optimal quantum adversary lower bounds for ordered search. In *Proc. 35th ICALP*, LNCS vol. 5125, pages 869–880, 2008, [arXiv:0708.3396 \[quant-ph\]](#).
- [FRPU94] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM J. Comput.*, 23(5):1001–1018, 1994. Earlier version in STOC’90.
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. 28th ACM STOC*, pages 212–219, 1996, [arXiv:quant-ph/9605043](#).
- [HLŠ07] Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proc. 39th ACM STOC*, pages 526–535, 2007, [arXiv:quant-ph/0611054](#).
- [HMW03] Peter Høyer, Michele Mosca, and Ronald de Wolf. Quantum search on bounded-error inputs. In *Proc. 30th ICALP*, pages 291–299, 2003, [arXiv:quant-ph/0304052](#). LNCS 2719.
- [HŠ05] Peter Høyer and Robert Špalek. Lower bounds on quantum query complexity. *EATCS Bulletin*, 87:78–103, October 2005, [arXiv:quant-ph/0509153](#).

- [Jor75] Camille Jordan. Essai sur la géométrie à n dimensions. *Bulletin de la S. M. F.*, 3:103–174, 1875.
- [KW93] Mauricio Karchmer and Avi Wigderson. On span programs. In *Proc. 8th IEEE Symp. Structure in Complexity Theory*, pages 102–111, 1993.
- [Lee09] Troy Lee. Composition upper bound for half-Boolean functions. unpublished, 2009.
- [MW05] Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122152, 2005, [arXiv:cs/0506068](#) [cs.CC].
- [NWZ09] Daniel Nagaj, Pawel Wocjan, and Yong Zhang. Fast amplification of QMA. *Quantum Inf. Comput.*, 9:1053–1068, 2009, [arXiv:0904.1549](#) [quant-ph].
- [Rei09a] Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. 2009, [arXiv:0904.2759](#).
- [Rei09b] Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. In *Proc. 50th IEEE FOCS*, pages 544–551, 2009.
- [Rei09c] Ben W. Reichardt. Span-program-based quantum algorithm for evaluating unbalanced formulas. 2009, [arXiv:0907.1622](#) [quant-ph].
- [Rei09d] Ben W. Reichardt. Faster quantum algorithm for evaluating game trees. 2009, [arXiv:0907.1623](#) [quant-ph].
- [Rei10] Ben W. Reichardt. Least span program witness size equals the general adversary lower bound on quantum query complexity. Technical Report TR10-075, Electronic Colloquium on Computational Complexity, <http://eccc.hpi-web.de>, 2010.
- [RŠ08] Ben W. Reichardt and Robert Špalek. Span-program-based quantum algorithm for evaluating formulas. In *Proc. 40th ACM STOC*, pages 103–112, 2008, [arXiv:0710.2630](#) [quant-ph].
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997, [arXiv:quant-ph/0508027](#). Earlier version in FOCS’94.
- [Sim97] Daniel R. Simon. On the power of quantum computation. *SIAM J. Computing*, 26(5):1474–1483, 1997. Earlier version in FOCS’94.
- [ŠS06] Robert Špalek and Mario Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2(1):1–18, 2006, [arXiv:quant-ph/0409116](#). Earlier version in ICALP’05.
- [Sze04] Mario Szegedy. Quantum speed-up of Markov chain based algorithms. In *Proc. 45th IEEE FOCS*, pages 32–41, 2004.

Institute for Quantum Computing, University of Waterloo
 E-mail address: breic@iqc.ca